Accurate Deep Representation Quantization with Gradient Snapping Layer for Similarity Search

Shicong Liu, Hongtao Lu

arXiv:1610.09645v1 [cs.CV] 30 Oct 2016

Abstract

Recent advance of large scale similarity search involves using deeply learned representations to improve the search accuracy and use vector quantization methods to increase the search speed. However, how to learn deep representations that strongly preserve similarities between data pairs and can be accurately quantized via vector quantization remains a challenging task. Existing methods simply leverage quantization loss and similarity loss, which result in unexpectedly biased back-propagating gradients and affect the search performances. To this end, we propose a novel gradient snapping layer (GSL) to directly regularize the back-propagating gradient towards a neighboring codeword, the generated gradients are un-biased for reducing similarity loss and also propel the learned representations to be accurately quantized. Joint deep representation and vector quantization learning can be easily performed by alternatively optimize the quantization codebook and the deep neural network. The proposed framework is compatible with various existing vector quantization approaches. Experimental results demonstrate that the proposed framework is effective, flexible and outperforms the state-of-the-art large scale similarity search methods.

Introduction

Approaches to large scale similarity search typically involves compressing a high-dimensional representation into a short code of a few bytes by hashing (Weiss, Torralba, and Fergus 2009; Gong and Lazebnik 2011) or vector quantization (Jegou, Douze, and Schmid 2011). Short code embedding greatly save space and accelerate the search speed. A trending research topic is to generate short codes that preserves semantic similarities, e.g., (Lin et al. 2014; Xiaojuan Wang 2016). Traditional short code generation methods require that the input data is firstly represented by some hand-crafted representations (e.g., GIST (Oliva and Torralba 2001) for images). Recently, learned representations with a deep neural network and a similarity loss is more preferable due to its stronger ability to distinguish semantic similarities(Chopra, Hadsell, and LeCun 2005; Wang et al. 2014).

In this paper we explore learning effective deep representations and short codes for fast and accurate similarity search. (Xia et al. 2014; Lai et al. 2015) propose to learn deep representations and mappings to binary hash codes with a deep convolutional neural network. However these deeply learned representations must exhibit certain clustering structures for effective short codes generation or severe performance loss could induce. How to learn representations that can be converted into short codes with minimal performance loss still remains a challenging task. Existing methods (Zhao et al. 2015; Zhu et al. 2016; Cao et al. 2016) add a quantization loss term to the similarity loss to regularize the output to be *quantizable*. However this methodology leads to difficulties in optimization, since they actually introduce a constant bias in the gradients computed from similarity loss for each sample. A joint representation and quantization codebook learning framework which is easy to be optimized and is compatible with various quantization methods is thus highly desirable.

Our main contributions are as follows: **(i)** We put forward a novel *gradient snapping layer* (GSL) to learn representations exhibiting required structures for vector quantization: Instead of directly regularizing the network output by adding a quantization loss, we regularize the back-propagating gradients by *snapping* the gradient computed from similarity loss toward an optimal neighbor codeword. By doing so we still reduce the quantization error and now the computed gradients are un-biased. **(ii)** GSL is yet another layer within existing deep learning framework and is compatible with any network structure. An illustrative configuration is presented in Fig.1. **(iii)** Our framework is compatible with any state-of-the-art vector quantization methods. GSL regularizes the output of its top layer to be *quantizable* by arbitrary vector quantization codebook. **(iv)** Experimental results on ILSVRC 2012 (Deng et al. 2009), CIFAR (Krizhevsky and Hinton 2009), MNIST (LeCun et al. 1998) and NUS-WIDE (Chua et al. 2009) show that our joint optimization framework significantly outperforms state-of-the-art methods by a large margin in term of search accuracy.

Related Works

Deep Representations for Similarity Search

Convolutional Neural Networks have shown strong performances in image classification tasks. Qualitative (Krizhevsky, Sutskever, and Hinton 2012) and quantitative (Babenko et al. 2014) evidences of their feasibility for image retrieval tasks

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

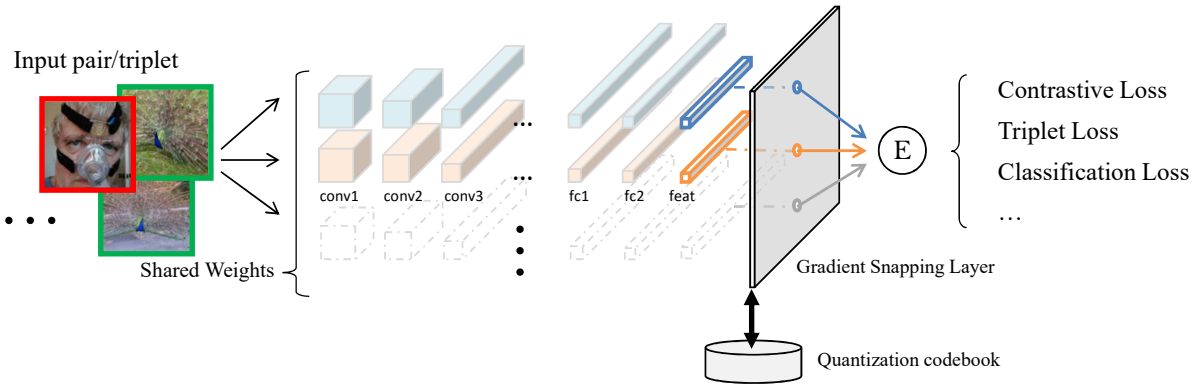

Figure 1: Illustrative network configuration with gradient snapping layer for deep representation quantization. We use a siamese network architecture, which takes several input images and computes network output with the same weights (also named shared weights). A gradient snapping layer associating a quantization codebook is inserted before the loss function, which could be any of the contrastive loss, triplet loss or other commonly used loss functions for neural network.

have also been observed in the former studies. Convolutional networks can generate effective deep representations within the *siamese architectures*(Chopra, Hadsell, and LeCun 2005), and are successful in fine-grained classification (Wang et al. 2014), unsupervised learning(Wang and Gupta 2015), and image retrieval(Xia et al. 2014; Lai et al. 2015; Zhao et al. 2015; Cao et al. 2016) tasks, etc. A Siamese architecture consists of several streams of layers sharing the same weights and uses a multi-wise loss function to train the network. A commonly used loss function for retrieval tasks is triplet loss(Lai et al. 2015), since it characterizes a relative ranking order and achieves satisfying performances(Kumar, Carneiro, and Reid 2015). Formally, given a sample $\mathbf{x}$, which is more similar to $\mathbf{x}^+$ rather than $\mathbf{x}^-$, the triplet loss is defined as :

$$E = max\{0, g + \|\hat{\mathbf{y}}_{\mathbf{x}} - \hat{\mathbf{y}}_{\mathbf{x}^+}\|_2 - \|\hat{\mathbf{y}}_{\mathbf{x}} - \hat{\mathbf{y}}_{\mathbf{x}_-}\|_2\} \quad (1)$$

, where $\hat{\mathbf{y}}_{\mathbf{x}}$, $\hat{\mathbf{y}}_{\mathbf{x}^-}$, $\hat{\mathbf{y}}_{\mathbf{x}^+}$ are the deep representations of $\mathbf{x}, \mathbf{x}_-, \mathbf{x}_+$, respectively. g is a gap parameter that regularizes the distance difference between "more similar" pair and "less similar" pair, it also prevents network from collapsing. Pair-wise loss (Cao et al. 2016) and classification loss are also used in recent literatures (Xiaojuan Wang 2016) due to their simplicity for optimization.

Compared to traditional hand-crafted visual features, deeply learned representations significantly improves the performances of image retrieval. However the deep representations obtained with convolutional neural networks are high dimensional (e.g., 4096-d for AlexNet (Krizhevsky, Sutskever, and Hinton 2012)), thus the representations must be effectively compressed for efficient retrieval tasks.

Vector Quantization for Nearest Neighbor Search

Vector quantization methods are promising for nearest neighbor retrieval tasks. They outperform the hashing competitors like ITQ(Gong and Lazebnik 2011), Spectral Hashing(Weiss, Torralba, and Fergus 2009), etc, by a large margin(Jegou, Douze, and Schmid 2011; Ge et al. 2013) on metrics like l_2. Vector quantization techniques involve quantizing a vector as a sum of certain pre-trained vectors and then compress this vector into a short encoding representation. Taking Product Quantization (PQ) (Jegou, Douze, and Schmid 2011) as an example, denote a *database* $\mathcal{X}$ as a set of N d-dimensional vectors to compress, PQ learns M *codebooks* $\mathcal{C}_1, \mathcal{C}_2, \cdots, \mathcal{C}_M$ for disjoint M subspaces, each of which is a list of K *codewords*: $\mathbf{c}_m(1)$, $\mathbf{c}_m(2)$, …, $\mathbf{c}_m(K) \in \mathbb{R}^{d/M}$. Then PQ uses a *mapping function* $i_m(\cdot) : \mathbb{R}^{d/M} \rightarrow [K]$[1] to encode a vector: $\mathbf{x} \mapsto (i_1(\mathbf{x}), \cdots, i_M(\mathbf{x}))$. A *Quantizer* q_m is defined as $q_m(\mathbf{x}) = \mathbf{c}_m(i(\mathbf{x}))$, meaning the m-th subspace of $\mathbf{x}$ is approximated as $q(\mathbf{x})$ for latter use. Product quantization minimizes the *quantization error*, which is defined as

$$E = \frac{1}{N} \sum_{\mathbf{x} \in \mathcal{X}} ||\mathbf{x} - [q_1(\mathbf{x}), \cdots, q_M(\mathbf{x})]^T||^2. \quad (2)$$

Minimizing Eqn.2 involves performing classical K-means clustering algorithm on each of the M subspaces. Afterwards, one can use *asymmetric distance computation* (ADC) (Jegou, Douze, and Schmid 2011) to greatly accelerate distance computation: Given a query vector $\mathbf{q}$, we first split it to M subspaces: $\mathbf{q} = [\mathbf{q}_1^\top, \cdots, \mathbf{q}_M^\top]^\top$. We then compute and store $\|\mathbf{q}_m - \mathbf{c}_m(k)\|^2$ for $k \in [K]$ and $m \in [M]$ in a look-up table. Finally, for each $\mathbf{x} \in \mathcal{X}$, the distance between $\mathbf{q}$ and $\mathbf{x}$ can be efficiently approximated using only $(M-1)$ floating-point additions and M table-lookups with the following equation:

$$\|\mathbf{q} - \mathbf{x}\|^2 \approx \sum_{m=1}^{M} \|\mathbf{q}_m - \mathbf{c}_m(i_m(\mathbf{x}))\|^2 \quad (3)$$

A number of improvements for PQ have been put forward, e.g., *Optimized product quantization* (OPQ) (Ge et al. 2013) *additive quantization* (AQ) (Babenko and Lempitsky 2014), and *generalized residual vector quantizaiton* (GRVQ) (Shicong Liu 2016), etc. Vector quantization methods also allow a number of efficient non-exhaustive search methods including Inverted File System(Jegou, Douze, and Schmid 2011), Inverted Multi-index(Babenko and Lempitsky 2012).

[1] $[K]$ denotes $\{1, 2, 3, \ldots, K\}$

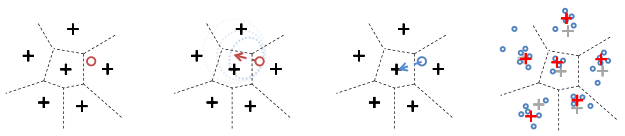

Figure 2: Illustrative demonstration of training with GSL. From left to right: 1) Red circle: raw representations; black crosses: codewords. 2) Compute the gradient from similarity loss, then enumerate neighboring codewords for gradient snapping. 3) Snap the gradient with Eqn.8. 4) Updated $\mathcal{C}$.

Recent studies on vector quantization focus on vector quantization for semantic similarities, for example, supervised quantization(Xiaojuan Wang 2016) applies a linear transformation that discriminates samples from different classes on the representations before applying composite quantization(Ting Zhang 2014). This approach is still upper-bounded by the input representations.

Joint Representation and Quantization Codebook Learning

To fully utilize deep representations for similarity search, the learned representation must strongly preserve similarities, and exhibit certain clustering structures to be *quantizable* for effective short code generation. This is a typical multi-task learning problem. Existing methods including deep quantization network (DQN) (Cao et al. 2016), deep hashing network (DHN) (Zhu et al. 2016) , simultaneous feature and hashing coding learning (SHFC)(Lai et al. 2015), etc, all propose to leverage between similarity loss $E(\cdot)$ and quantization loss $Q(\mathbf{y}) = \|\mathbf{y} - q(\mathbf{y})\|^2$, and derive the final loss function[2]:

$$E'(\mathbf{y}) = E(\mathbf{y}) + \lambda Q(\mathbf{y}) \tag{4}$$

However, the above loss function is actually not suitable for joint representation and quantization codebook learning. Consider the following:

$$\begin{aligned}\frac{\partial E'(\mathbf{y})}{\partial \mathbf{y}} &= \frac{\partial E(\mathbf{y})}{\partial \mathbf{y}} + \lambda \frac{\partial Q(\mathbf{y})}{\partial \mathbf{y}} \\ &= \frac{\partial E(\mathbf{y})}{\partial \mathbf{y}} + 2\lambda(\mathbf{y} - q(\mathbf{y})) \\ &= \frac{\partial E(\mathbf{y})}{\partial \mathbf{y}} + 2\lambda\epsilon_{\mathbf{y}}\end{aligned} \tag{5}$$

where $\epsilon_{\mathbf{y}}$ denotes the quantization residual vector of $\mathbf{y}$. Thus existing methods can be considered as simply adding a bias term specific to each sample. This is not desirable since we want to find a gradient that both reduce the similarity loss and the quantization loss, the biased similarity loss distorts the gradients computed from similarity loss and eventually leads to sub-optimal representations.

Gradient Snapping Layer

We propose *gradient snapping layer* (GSL) to jointly learn representations and to enforce the learned representations to exhibit cluster structure. We first start from the following loss function leveraging similarity loss and an adaptive quantization loss:

$$\begin{aligned}&E'(\mathbf{y}) = \lambda_1 E(\mathbf{y}) + \lambda_2 \|\mathbf{y} - \mathbf{c}\|^2, \\ &\text{s.t. } \arg\max_{\mathbf{c}\in\mathcal{C}} (\frac{\partial E(\mathbf{y})}{\partial \mathbf{y}})^T \frac{\partial E'(\mathbf{y})}{\partial \mathbf{y}} \cdot f(\|\mathbf{c} - \mathbf{y}\|^2)\end{aligned} \tag{6}$$

$f(\cdot) \in \mathbb{R}^+$ is a function that leverages between different neighboring codewords of $\mathbf{y}$. A neighboring codeword is more preferred compared to distant codeword. We may choose a Gaussian function $f(d) = e^{-d^2/\sigma}$ where σ is the mean l_2 distance between $\mathbf{y}$ and the neighboring codewords. It's easy to see that the above loss function is unbiased by introducing a small variance to the quantization loss term. Eqn.4 is a special case when we choose $\mathbf{c} = q(\mathbf{y})$.

Based on Eqn.6, we can learn *quantizable* and similarity preserving representations by GSL by inserting it between the representation layer and the similarity loss layer. It works like snapping the back-propagating gradients towards a codeword: On the forward-pass phase, it directly outputs the representation vector $\mathbf{y}$. On the backward-pass phase, it takes gradient $\frac{\partial E}{\partial \mathbf{y}}$ from the similarity loss layer, and solves Eqn.6 by enumerating a number of nearest codewords $\mathbf{c} \in \mathcal{C}$ to $\mathbf{y}$ and minimizing the following:

$$\begin{aligned}&\arg\max_{\mathbf{c}} (\frac{\partial E}{\partial \mathbf{y}})^T \cdot \Delta\mathbf{c}, \\ &\text{where} \quad \Delta\mathbf{c} = f(\|\mathbf{c} - \mathbf{y}\|^2) \frac{(\mathbf{c} - \mathbf{y})}{\|\mathbf{c} - \mathbf{y}\|}\end{aligned} \tag{7}$$

Then we compute and back-propagates the snapped gradient $\Delta\mathbf{y}$ to the representation layer. In this paper we choose a simple yet effective strategy that leverages the original gradient and its projection towards codeword $\mathbf{c}$, which is given below:

$$\begin{aligned}\Delta\mathbf{y} &= \lambda_1 \frac{\partial E}{\partial \mathbf{y}} + \lambda_2 \Delta\mathbf{c}, \\ \text{where } \lambda_1 &= (1 - \frac{((\frac{\partial E}{\partial \mathbf{y}})^T \Delta\mathbf{c})^2}{\|\Delta\mathbf{c}\|^2 \|\frac{\partial E}{\partial \mathbf{y}}\|}) \lambda, \\ \lambda_2 &= \frac{(\frac{\partial E}{\partial \mathbf{y}})^T \Delta\mathbf{c}}{\|\Delta\mathbf{c}\|}\end{aligned} \tag{8}$$

, where λ_2 is the projected gradients on $\mathbf{c}$, λ_1 computes the residual of original gradients by projecting $\lambda_2\Delta\mathbf{c}$ back and further scaled by λ.

Simultaneous to representation training, $\mathcal{C}$ should also be updated to fit the new deep representations. A simple method is to gather the new representations and update $\mathcal{C}$ every a number of iterations. Take PQ as an example, one can use sequential k-means clustering on each subspace streamed with network output to update the codebooks. We present an illustrative explanation of how our gradient snapping layer work in Fig.2 for better understanding.

Note GSL can be seen as yet another type of layer within deep learning framework, and the main use is to regularize the learned representations to be *quantizable*. Thus it's naturally compatible with any loss function (e.g., triplet loss,

[2]The methodology is the same though they could adopt different similarity/quantization loss function.

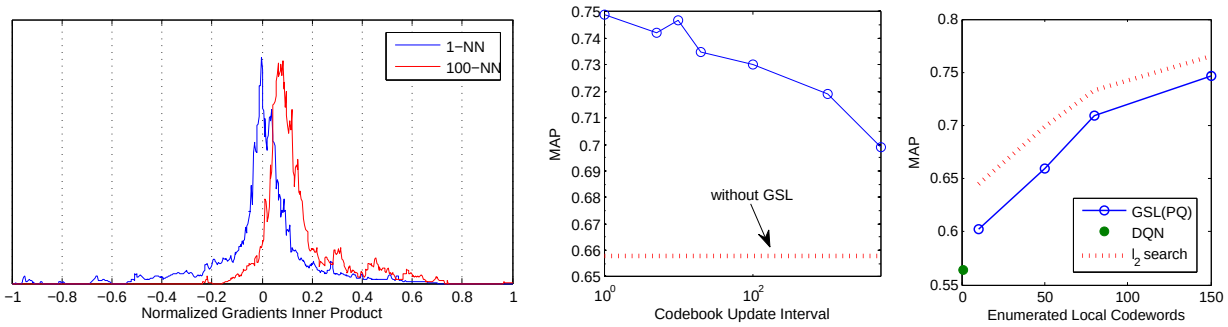

(a) (b) (c)

Figure 3: Empirical analysis on CIFAR10, using triplet loss, and associate with 32-bit product quantization codebook. a) Comparison between output regularization and gradient snapping. We normalize $\partial E/\partial \mathbf{y}$ and $\Delta \mathbf{c}$, then draw the histogram of their inner products. b) The impact of codebook update interval on retrieval performance measured in MAP (mean average precision), with 150 enumerated codewords. c) How the number of enumerated neighboring codewords impacts retrieval performance, with codebook update interval 1.

Table 1: Image retrieval performances in MAP of 32-bit code obtained with different representations on CIFAR-10. Value in brackets are reported from (Lai et al. 2015; Shen et al. 2015); values with * are trained on all 59K training samples.

	l_2	PQ	FashHash	SDH	ITQ	AQ	OPQ	KSH
GIST	*0.177*	0.170	0.372 (~0.37*)	**0.400** (~0.43*)	0.174 (0.172)	0.175	0.178	(0.346)
AlexNet 4096-d	*0.313*	0.264	**0.603**	0.586	0.260	0.233	0.292	(0.558)
feat, no GSL	*0.779*	0.658	0.646	0.619	0.582	0.669	**0.683**	-
feat+GSL(PQ)	*0.765*	**0.747**	0.641	0.624	0.603	0.691	-	-
feat+GSL(AQ)	*0.751*	0.639	0.638	0.615	0.590	**0.712**	-	-

contrastive loss) or any network structures. We present an typical configuration in Figure 1. The learned representations are then quantized, stored, queried linearly or with systems like IVFADC and inverted-multi index. Readers may refer to (Jegou, Douze, and Schmid 2011; Babenko and Lempitsky 2012; Kalantidis and Avrithis 2014) for details.

Experiments and Discussions

We conducted extensive experiments on the following commonly used benchmark datasets to validate the effectiveness of gradient snapping layer (GSL):

NUS-WIDE(Chua et al. 2009), a commonly used web image dataset for image retrieval. Each image is associated with multiple semantic labels. We follow the settings in (Lai et al. 2015; Cao et al. 2016), and use the images associated with the 21 most frequent labels and resize the images into 256×256. [3]

CIFAR-10(Krizhevsky and Hinton 2009), a dataset containing 60,000 32×32 tiny images in 10 classes commonly used for image classification. It's also a popular benchmark dataset for retrieval tasks. The dataset contains 512-d GIST feature vectors for all images.

[3]We gather and use 139,633 images for experiment due to a number of original images are no longer available.

MNIST(LeCun et al. 1998), a dataset of 70,000 28×28 greyscale handwritten digits from 0 to 9. Each image is represented by raw pixel values, resulting a 784-d vector.

ILSVRC2012(Deng et al. 2009), a large dataset containing over 1.2 million images of 1,000 categories. It has a large training set of 1,281,167 images, around 1K images per category. We use the provided validation images as the query set, it contains 50,000 images.

Following previous works of (Xia et al. 2014) (Lai et al. 2015)(Cao et al. 2016), for MINST, CIFAR-10 and NUS-WIDE, we sample 100 images per class as the test query set, and 500 images per class as training set. We use all training data in ILSVRC2012 to train the network as in (Xiaojuan Wang 2016) and use the validation dataset as query.

We evaluate image retrieval performance in term of mean average precision (MAP). We use hamming ranking for hashing based methods and asymmetric distance computation for vector quantization to retrieve the image rankings of the dataset. We limit the length of retrieved images to 1500 for ILSVRC2012 as in (Xiaojuan Wang 2016), and 5000 for NUS-WIDE as in (Xia et al. 2014) etc.

Network Settings

We implement GSL on the open-source Caffe deep learning framework. We employ the AlexNet architecture(Krizhevsky, Sutskever, and Hinton 2012) and use pre-trained weights

Table 2: Comparison between different simultaneous representation and short code learning schemes. *CNNH is under a two-step hashing framework. Other methods perform simultaneous representation and short code learning in an end-to-end framework.

	Loss function	Regularization	Search
DQN(Cao et al. 2016)	Pair-wise cosine	Product quantization	ADC
DHN(Zhu et al. 2016)	Pair-wise logistic	l_1 quantization	Hamming
CNNH(Xia et al. 2014)*	Classification	No	Hamming
DRSCH(Zhang et al. 2015)	Triplet & pair-wise	No	Weighted Hamming
DSFH(Zhao et al. 2015)	List-wise	Zero-mean	Hamming
SHFC(Lai et al. 2015)	Triplet	No	Hamming
Ours	Any	Gradient snapping	ADC

for convolutional layers *conv1-conv5, fc6, fc7*. We stack an additional 192-d inner product layer *feat*, GSL layer and triplet-loss layer on *fc7*. We choose triplet loss due to it's typical and commonly used. It also preserve similarity in l_2 which is compatible with most existing vector quantization methods. We use the activation at *feat* as deep representation. We employed the triplet selection mechanism described in (Wang et al. 2014) for faster convergence. Small projection transformations are applied to the input images for data augmentation, we choose $\lambda = 0.036$. We'll make our implementation publicly available to allow further researches.

Empirical Analysis

We perform experiment on CIFAR-10 to investigate the performances of image retrieval using PQ(Jegou, Douze, and Schmid 2011), AQ(Babenko and Lempitsky 2014), FashHash(Lin et al. 2014), SDQ(Shen et al. 2015), ITQ(Gong and Lazebnik 2011), KSH(Liu et al. 2012) and using exhaustive l_2 search with 512-d GIST representation provided by the dataset, 4096-d activation on layer *fc7* of pretrained AlexNet on ImageNet, and AlexNet fine-tuned with GSL associated with PQ/AQ-32 bit codebook for representation compressing. We retrieve 150 nearest codewords and compute the snapped gradient. We choose 3000 anchors for SDH, 4 as tree depth for FashHash, $M = 4, K = 256$ for PQ/AQ. The results are shown in Table 1.

Effectiveness of end-to-end learning: The experimental result again validate the importance of end-to-end learning of representation and short-codes as addressed in (Zhao et al. 2015; Zhang et al. 2015; Zhu et al. 2016). A fine-tuned representation allow further performance improvements.

Regularization effectiveness: Experiments on CIFAR-10 demonstrate GSL has a significant effect on regularizing the deep representations to be quantized accurately. Specifically, compared to representations without regularization, we observed a 8.9% performance improvement for PQ and 4.3% for AQ. This regularization doesn't induce a heavy penalty on the performance of l_2 exhaustive search, which can be seen as the performance upper bound for vector quantization methods.

Vector quantization methods for GSL: The choice of different vector quantization method has a small influence on exhaustive l_2 search. PQ performs better since it can enumerate neighboring codewords deterministically, so we use PQ for the following experiments. One should always use the same vector quantization method for training and retrieval, which is obvious. Further tests on CIFAR-10 with GSL(PQ)-32bit show that: 1> A higher update frequency of $\mathcal{C}$ improves performances as shown in Fig.3(b). 2> Retrieval performances improves when more neighboring codewords are enumerated during training as shown in Fig.3(c), however it also requires more computation resources.

Gradient Snapping vs Output Regularization: Existing simultaneous representation and hash coding learning (Zhao et al. 2015; Zhang et al. 2015; Zhu et al. 2016) or quantization codebook learning (Cao et al. 2016) methods add a penalty term for quantization error for output regularization. This can be seen as a special case of our proposed method when we reduce the number of enumerated neighboring codewords to one, which is then equivalent to adding a constant bias for each sample. In fact, existing vector quantization methods, e.g., (Jegou, Douze, and Schmid 2011; Ge et al. 2013; Shicong Liu 2016; Babenko and Lempitsky 2014) generate codewords combinatorially, and codewords are quite dense in the local space of a representation vector(Jegou, Douze, and Schmid 2011), thus we can optimize the quantization loss w.r.t. any neighboring codewords. Then the back-propagated gradients can be more conform to the gradients from similarity loss. This makes the learned representations better preserving similarities. We normalize the both gradients and compute the inner-product for each sample on each iteration and report it in Fig.3(a), our GSL achieves significantly better conformity between the two gradients.

Note that in extreme cases $(\partial E/\partial \mathbf{y})^T \cdot \Delta \mathbf{c} < 0$ for all neighboring codewords, we then set $\lambda_1 = \lambda$ and $\lambda_2 = 0$ to reject snapping. A better handling of this case might yield better performances.

Performance Comparison

We briefly list the state-of-the-arts simultaneous representation and short code learning methods and the comparison between them in Table2. Since we adopt the same experimental settings, we directly report the experimental result from the corresponding papers. The image retrieval performances measured in MAP on CIFAR10, MNIST, NUS-WIDE are presented in Table 3. Our method out-performs state-of-the-art deep hashing/quantization methods by large margin. We also present retrieval precision w.r.t. top returned samples curves for CIFAR-10 and NUS-WIDE with 48-bit short code obtained with different methods in Fig.4(a) and Fig.4(b).

Table 3: Mean Average Precision (MAP) for Different Number of Bits on three benchmark datasets.

	CIFAR				MNIST				NUS-WIDE			
Method	12-bit	24-bit	32-bit	48-bit	12-bit	24-bit	32-bit	48-bit	12-bit	24-bit	32-bit	48-bit
CNNH*	0.484	0.476	0.472	0.489	0.969	**0.975**	0.971	0.975	0.617	0.663	0.657	0.688
SHFC	0.552	0.566	0.558	0.581		N/A			0.674	0.697	0.713	0.715
DQN	0.554	0.558	0.564	0.580		N/A			0.768	0.776	0.783	0.792
DHN	0.555	0.594	0.603	0.621		N/A			0.708	0.735	0.748	0.758
DRSCH	-	0.622	0.629	0.631	-	0.974	0.979	0.971		N/A†		
DSCH	-	0.611	0.617	0.618	-	0.967	0.972	0.975		N/A†		
Ours	-	**0.692**	**0.747**	**0.752**	-	0.973	**0.980**	**0.981**	-	**0.812**	**0.839**	**0.849**

*An enhanced variant of CNNH proposed in (Lai et al. 2015)
† (Zhang et al. 2015) uses a different experimental setting.

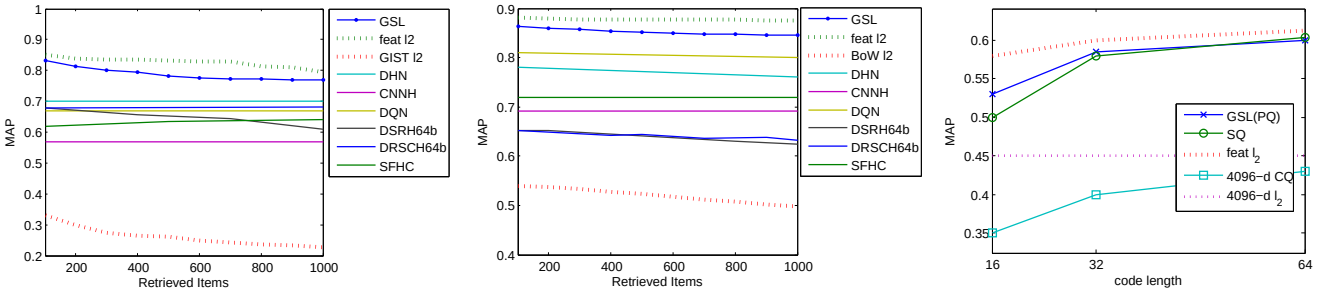

(a) CIFAR10 (b) NUS-WIDE (c) ILSVRC2012

Figure 4: Image retrieval performance comparison.

Please note that previous deep representation learning schemes using triplet loss for hashing, e.g.(Lai et al. 2015; Zhang et al. 2015) are outperformed by schemes using pair-wise loss, e.g., (Cao et al. 2016; Zhu et al. 2016). This doesn't mean triplet loss is inferior to pair-wise loss. In fact, it has been testified that triplet loss works better than pair-wise loss on various tasks including classification(Wang et al. 2014; Norouzi, Fleet, and Salakhutdinov 2012) and face retrieval(Schroff, Kalenichenko, and Philbin 2015). The main reason is that the quantization regularization introduced in previous triplet loss based method is too weak compared to recent pair-wise loss based methods. By introducing better regularization method, our method achieves significantly better result with triplet loss..

We also compared large scale image retrieval on ILSVRC2012 in Figure 4(c). We report the MAP results from (Xiaojuan Wang 2016) for CQ and SQ. SQ, CQ uses the 4096-d feature extracted from layer *fc7* of AlexNet pre-trained on ILSVRC2012. Thus SQ can be seen as adding a "final inner product layer" with output regularization on AlexNet, however SQ doesn't allow the weights of the network to be updated. Our end-to-end representation and short code learning improves the retrieval performances.

Conclusion

In this paper we propose to regularize gradients on deep convolutional neural network to obtain deep representations exhibiting structures compatible for accurate vector quantization. A gradient snapping layer is discussed detailedly to this end. We validate our approach by training a siamese network architecture based on AlexNet with triplet loss functions to learn representations for similarity search. The experimental results show our approach enforces quantization regularization while learning strong representation, and outperforms the state-of-the-art similarity search methods.

References

Babenko, A., and Lempitsky, V. 2012. The inverted multi-index. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, 3069–3076. IEEE.

Babenko, A., and Lempitsky, V. 2014. Additive quantization for extreme vector compression. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, 931–938. IEEE.

Babenko, A.; Slesarev, A.; Chigorin, A.; and Lempitsky, V. 2014. Neural codes for image retrieval. In *Computer Vision–ECCV 2014*. Springer. 584–599.

Cao, Y.; Long, M.; Wang, J.; Zhu, H.; and Wen, Q. 2016. Deep quantization network for efficient image retrieval. In *Thirtieth AAAI Conference on Artificial Intelligence*.

Chopra, S.; Hadsell, R.; and LeCun, Y. 2005. Learning a similarity metric discriminatively, with application to face verification. In *Computer Vision and Pattern Recognition,*

2005. CVPR 2005. IEEE Computer Society Conference on, volume 1, 539–546. IEEE.

Chua, T.-S.; Tang, J.; Hong, R.; Li, H.; Luo, Z.; and Zheng, Y. 2009. Nus-wide: a real-world web image database from national university of singapore. In *Proceedings of the ACM international conference on image and video retrieval*, 48. ACM.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, 248–255. IEEE.

Ge, T.; He, K.; Ke, Q.; and Sun, J. 2013. Optimized product quantization for approximate nearest neighbor search. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, 2946–2953. IEEE.

Gong, Y., and Lazebnik, S. 2011. Iterative quantization: A procrustean approach to learning binary codes. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, 817–824. IEEE.

Jegou, H.; Douze, M.; and Schmid, C. 2011. Product quantization for nearest neighbor search. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 33(1):117–128.

Kalantidis, Y., and Avrithis, Y. 2014. Locally optimized product quantization for approximate nearest neighbor search. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, 2329–2336. IEEE.

Krizhevsky, A., and Hinton, G. 2009. Learning multiple layers of features from tiny images.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.

Kumar, B.; Carneiro, G.; and Reid, I. 2015. Learning local image descriptors with deep siamese and triplet convolutional networks by minimising global loss functions. *arXiv preprint arXiv:1512.09272*.

Lai, H.; Pan, Y.; Liu, Y.; and Yan, S. 2015. Simultaneous feature learning and hash coding with deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3270–3278.

LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324.

Lin, G.; Shen, C.; Shi, Q.; Hengel, A.; and Suter, D. 2014. Fast supervised hashing with decision trees for high-dimensional data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1963–1970.

Liu, W.; Wang, J.; Ji, R.; Jiang, Y.-G.; and Chang, S.-F. 2012. Supervised hashing with kernels. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, 2074–2081. IEEE.

Norouzi, M.; Fleet, D. J.; and Salakhutdinov, R. R. 2012. Hamming distance metric learning. In *Advances in neural information processing systems*, 1061–1069.

Oliva, A., and Torralba, A. 2001. Modeling the shape of the scene: A holistic representation of the spatial envelope. *International journal of computer vision* 42(3):145–175.

Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 815–823.

Shen, F.; Shen, C.; Liu, W.; and Tao Shen, H. 2015. Supervised discrete hashing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 37–45.

Shicong Liu, Junru Shao, H. L. 2016. Generalized residual vector quantization for approximate nearest neighbor search. In *International Conference on Multimedia and Expo, 2016. ICME 2016, to appear*.

Ting Zhang, Chao Du, J. W. 2014. Composite quantization for approximate nearest neighbor search. *Journal of Machine Learning Research: Workshop and Conference Proceedings* 32(1):838–846.

Wang, X., and Gupta, A. 2015. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE International Conference on Computer Vision*, 2794–2802.

Wang, J.; Song, Y.; Leung, T.; Rosenberg, C.; Wang, J.; Philbin, J.; Chen, B.; and Wu, Y. 2014. Learning fine-grained image similarity with deep ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1386–1393.

Weiss, Y.; Torralba, A.; and Fergus, R. 2009. Spectral hashing. In *Advances in neural information processing systems*, 1753–1760.

Xia, R.; Pan, Y.; Lai, H.; Liu, C.; and Yan, S. 2014. Supervised hashing for image retrieval via image representation learning. In *AAAI*, volume 1, 2.

Xiaojuan Wang, Ting Zhang, G.-J. Q. J. T. J. W. 2016. Supervised quantization for similarity search. In *Computer Vision and Pattern Recognition, 2016. CVPR 2016, to appear*.

Zhang, R.; Lin, L.; Zhang, R.; Zuo, W.; and Zhang, L. 2015. Bit-scalable deep hashing with regularized similarity learning for image retrieval and person re-identification. *Image Processing, IEEE Transactions on* 24(12):4766–4779.

Zhao, F.; Huang, Y.; Wang, L.; and Tan, T. 2015. Deep semantic ranking based hashing for multi-label image retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1556–1564.

Zhu, H.; Long, M.; Wang, J.; and Cao, Y. 2016. Deep hashing network for efficient similarity retrieval. In *Thirtieth AAAI Conference on Artificial Intelligence*.